DOCUMENT RESUME

ED 112 684 FL 007 164

AUTHOR Wilks, Yorick
TITLE Parsing English. Course Notes for a Tutorial on Computational Semantics, March 17-22, 1975.
PUB DATE 25 Jan 75
NOTE 68p.; Paper given at the Institute for Semantic and Cognitive Studies (Castagnola, Switzerland, March 17-22, 1975); For related document, see FL 007 167
AVAILABLE FROM Institute for Semantic and Cognitive Studies, Villa Heleneum, 6976 Castagnola, Switzerland ($10.00 for complete volume)

EDRS PRICE MF-$0.76 HC-$3.32 Plus Postage
DESCRIPTORS Artificial Intelligence; *Computational Linguistics; *Computer Programs; Form Classes (Languages); Grammar; Linguistic Theory; Programing; Programing Languages; *Semantics; *Structural Analysis; *Syntax; Transformation Generative Grammar
IDENTIFIERS *Parsing

ABSTRACT

The course in parsing English is essentially a survey and comparison of several of the principal systems used for understanding natural language. The basic procedure of parsing is described. The discussion of the principal systems is based on the idea that "meaning is procedures," that is, that the procedures of application give a parsed structure its significance. Natural language systems should be content- rather than structure-motivated, i.e. they should be concerned with linguistic problems revealed by parsing rather than with the relation of the proposed structure of the system to the structures of other systems. Within this framework, Winograd's understanding system, SHRDLU, is described and discussed, as are the second generation systems of Simmons, Schank, Colby and Wilks. A subsequent discussion compares all these systems. Concluding remarks outline immediate problems, including the need for a good memory model and the use of texts, rather than individual example sentences, for investigation. (CLK)

SYNOPSIS/NOTES FOR COURSE PARSING ENGLISH

Yorick Wilks - Jan 25 1975

1. Introduction

This course will be, in essence, a survey and comparison of several of the principal systems for understanding natural language now under discussion in the Artificial Intelligence (AI) community. The dimensions of comparison chosen will reflect the presuppositions and prejudices of the lecturer, and I shall try to make some of these clear right at the beginning.

Given any proposed structure for natural language whatever, we may distinguish between the structure itself, and its correlation with pieces of language. Among such correlations we may distinguish between parsings and assignments. By "parsing" here I mean the provision of definite procedures of application, and by assignment I mean the provision of no more than a list of correspondences, between chunks of language and formal structures. Logicians often tell us that, say, the structure of the sentence "John loves his wife's sister" is of the form Ex.Ey.Ez:x=J. y=W(J). z=S(W(J)).L(x z),but they rarely, if ever, provide a parsing of that structure onto the sentence, for the procedure is considered obvious. That is what I mean by an assignment, or the provision of a list of correspondences between sentences and structures.

At this point let me try and say a little what a parsing is rather than what it is not. Let us consider the structure specified by one of the simplest grammars: a content-free phrase structure grammar. This might contain rules like the following (an example of Winograd's):

S → NP VP
NP → Determiner Noun
VP → Verb NP
Determiner → a, the
Noun → rabbit, cucumber
Verb → nibbled

Sentence: The rabbit nibbled a cucumber.

This sentence "The rabbit nibbled the cucumber" is considered to be produced by the listed rewrite rules of this grammar. The grammar is thought of a implicitly defining a procedure; starting with the symbolS, and re-writing it as the right hand side of the rule containing it, and then rewriting those symbols by other rules (or the same rule can be used again if it is possible to do so), as long as only one symbol is rewritten each time a rule is applied. The production of this sentence is usually represented by a phrase-structure tree as follows (also called a phrase-marker) where every production of branches from a node corresponds to the application of one of these rewrite rules. As just described, eventually symbols are written that are actually words of English, and the process can go no further when every node is an English word.

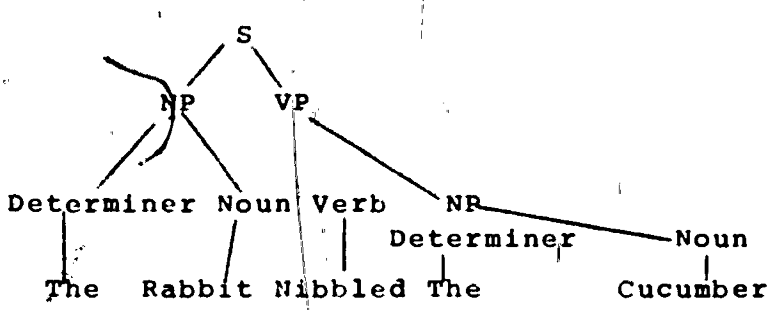

However, this process of idealised sentence generation is not a parsing process for that is always from the sentence and some such battery of rules to a structure like that of the diagram. The structure obtained is then called a parsing because it reflects certain syntactic relationships that we may think of as holding between the words of the sentence. Thus in the tree structure, "the" has been tied to "rabbit" by the application of the rule NP $\longrightarrow$ Determiner Noun and that relation of "the" to "rabbit" is the dependence of a determiner on the noun it determines.

For the purpose of parsing we think of each word in the sentence to be parsed as attached to one or more grammatical categories: in this case as we see from the rules "rabbit" is attached to a single category "Noun". There are two ways of doing the parsing: top-down and bottom-up. Bottom-up is the more straightforward way, and is illustrated by the next figure (all figures in this part are due to Winograd). The words of the sentence are listed and starting from the left of the sentence, we attempt to replace each one by its category, and then to rewrite successively pairs of category symbols by reversing the grammar's rewrite rules, until we reach the final sentence symbol S. The lines of the derivation (the tree above upside down, in fact) can then be considered as the parsing.

The rabbit nibbled a cucumber.
Determiner rabbit nibbled a cucumber.
Determiner Noun nibbled a cucumber.
NP nibbled a cucumber
NP Verb a cucumber
NP Verb Determiner cucumber.
NP Verb Determiner Noun
NP Verb NP
NP VP
S

Top down parsing starts, not with the words and their categories, but with generations as described initially form the top line of the grammar starting with symbol S. It continues to generate sequences of categories until it finds one that will match directly onto the leftmost word of the sentence, "The". This procedure continues,

as illustrated by the next diagram, until the last word at the right is reached and then once again the lines of the procedure are the parsing. The two derivations above (bottom up and top down) are formally equivalent.

Looking for: S : |the rabbit nibbled a cucumber.
Looking for: NP | VP : |the rabbit nibbled a cucumber.
Looking for: Determiner | Noun VP : |the rabbit nibbled a cucumber.
Looking for: Noun | VP : the | rabbit nibbled a cucumber.
Looking for: VP : the rabbit | nibbled a cucumber.
Looking for: Verb | NP : the rabbit | nibbled a cucumber.
Looking for: NP : the rabbit nibbled | a cucumber.
Looking for: Determiner | Noun : the rabbit nibbled | a cucumber.
Looking for: Noun : the rabbit nibbled a | cucumber.
Looking for: : the rabbit nibbled a cucumber|.

Several things should be noted here: parses do not <u>have</u> to be left to right, they could easily be the other way round. The structure is rarely as trivial as this one, and usually a word will have several possible categories assigned to it in the dictionary and part of the job of the parsing process is to find out which of those roles the word has in any particular sentence. So "run" has a Verb category in "I run for a bus" but a Noun category in "I built a new chicken run", (a run is a cage for chickens). So, if we used either of the above processes to parse "I run for a bus", we would find that we produce the correct structure:

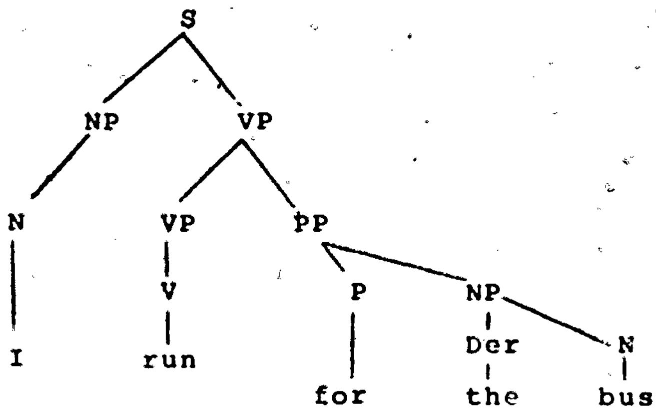

but never a structure where "run" had its N(Noun) category, because there is no sequence of proper grammar rule applications that could produce a tree for that sentence with the node above "run" labelled N. ~~In~~ particular, there is a rule:

VP ⟶ V (and then V ⟶ run)

but no rule:

VP ⟶ N (and then N ⟶ run)

So a "reading" of "run" is excluded bz the parsing of this sentence.

At this point, another important pair of technical terms enter: <u>breadth-first</u> and <u>depth-first</u> parsings. Breadth first is the parallel treatment of all possible alternative structures at a given time, none is given any precedence. They are all developed until one or more reach the success symbol S (if the parsing is bottom-up). In depth-first parses the alternative structures are treates sequentially, and if one combination is no longer a possible success (because, in this grammar no more rewrites of a string are possible even though the parsing has not been completed), then the system must back-track, which means going back and continuing one of the alternative structures that had not yet been tested.

In the case of transformational structures in particular, a sentence is often parsed so as to produce a structure quite different from the original form of the sentence. So, to take a classic example, a passive sentence like "I was given the money last week" might be given a parsed structure representing it whose form was actually that of "Someone gave me the money last week", as regards the order of its principal constituents. So we may say that a parsing may produce a structure quite different from the superficial form of the input sentence. Grammatical parsings programs of this sort were a growth industry in the sixties, but I do not want to go into them in any detail here beyond the introduction of the basic notions. That is because one of the assumptions of this course is that grammatical (or syntactic) parsing of the sort described is not fundamental, and that is need not be even a preliminary to assigning a useful meaning structure to sentences. However, I shall argue that only indirectly, and much later on. I shall go into some detail describing the syntactic parser behind Terry Winograd's program. Later, when I come to my own system (and that of Chris Riesbeck) I shall describe parsing procedures (that is to say procedural ascription of structures to sentences) where the structures are not syntactic at all, but semantic, and at that point I shall deal with some of the difficulties of manipulating syntactic information by quite other means.

Let us return for a moment to what I called <u>assignments</u>. An intermediate case between parsings and assignments would be when a sentence was proferred as its own representation, so that its structure was no more than what

a reader of the sentence would intuitively assign to it, were he asked to do so. That is to say, if any writer were to offer us "John loves his wife's sister" as the representation of "John loves his wife's sister", we would know something odd was going on. Things are rarely as bad as this, but they do approach this situation, one might say, asymptotically. So for example, in any AI language program in which the representation for "John is at the station" is given as (AT JOHN STATION), and this is quite a common move as you will see from other courses in this tutorial, then we are being given a structure which is only trivially different from that of the sentence itself, and the reader of such a structure must "parse" it intuitively by some direct matching such as :

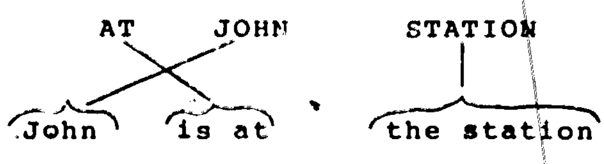

I would claim that such a process is also not a serious parsing, but rather what I would call a projection, in that the sentences chosen as examples are essentially projections of the representation itself,and moreover, and this is the important point, the representation is trivial and can only be understood from our knowledge of the original sentence. If the representation and sentence were in a foreign language, of which one knew nothing, one would soon see that without explicit rules of interpretation for the representation (which we are not normally given) one could not understand it.

A full, complex, and explicit representation can also be projected in some sense, of course, but at the moment I shall restrict the term to simplistic representations.

The effect of both assignments and projections is to make parsing seem either unnecessary, or at most a detachable and ultimately dispensible "front end" to some other system. It will become clear that I think this assumption to be radically mistaken, in that, as I am using "parsing", it is not only a front end, and not only a test of the value and efficacy of the representation parsed, for it also gives meaning to the structure.

My view is a version of the "meaning is procedures" attitude: that the procedures of its application give a parsed structure significance, and that without the procedure or the demonstrated possibility of them, the structure proposed is essentially meaningless, and the plaything of formalists.

At this point, I shall consider, in passing, a modern movement in linguistics, Generative Semantics, which has got into a position in which, given its theoretical presuppositions, the structures it postulates cannot, demonstrably be parsed onto sentences. I shall argue this in detail and point out that the position is not made any less serious by the fact that most of the practitioners of Generative Semantics have not noticed it.

Let me make a second distinction at this point that will be important in the examination of systems that follows. It is between AI language systems that are primarily content-motivated and those which are primarily structurally-motivated I shall not claim that systems are of one sort or the other, since there is as always, a continuum in these matters, but will ask readers to keep the distinction in mind during the description and discussion of systems.

By structurally-motivated systems I mean those which, whatever their natural language capacities, are set up in order to solve problems that are essentially non-linguistic; and are concerned with the relation of the proposed representation of structure to other structures. To put it another way, structurally-motivated systems are not concerned with the problems posed by the extraction of the content of input natural language, since the nature and purpose of the representation adopted is aimed at the problem rather than at language.

A paradigm of this would be Bobrow's early "Natural Language System" (Bobrow '68) that solved elementary algebra problems. It did indeed take in a form of restricted English, but it, in my view, was essentially a system to solve algebra problems. In contrast to this, one cannot even imagine a machine translation system that accepted and produced structures: for the extraction of content is the problem the whole system is set up to solve.

It will be clear that my distinctions are evaluative in nature, and indeed one of the purposes of this course will be to argue in an indirect way that natural language systems, 1) should be parsed rather than assigned or projected and 2) should be content-rather than structurally-motivated. I certainly cannot prove, however, that advance with "The problem of natural language" will come through content-rather than structurally-motivated systems. NOtice too, here that the two distinctions are not necessarily connected: Winograd's system, for example, that I shall discuss in some detail later, is certainly a parsed system, yet is it arguably structurally-motivated.

What do I mean by "the problem of natural language"? I mean the problem that has dogged mechanical language analysis since its inception in the Fifties, and which centres round the problem of systematic ambiguity. Let me give a brief and tendentious historical sketch at this point: I take it that the "machine translation era" of the Fifties and early Sixties was a bold attempt to do an enormous taks: extract the content of natural language utterances. It failed, and for three reasons: word sense ambiguity or polysemy, case ambiguitv (the ambiguity of prepositions) and referential ambiguity (the ambiguity of pronouns). It failed to develop and implement structures to deal with these features of even non-metaphorical natural language. Generative linguistics was born partly in response to this failure: it has provided complex structures, but has lost the sense of any definite task to be performed, and the most recent sign of this failure of nerve has been a retreat from any relation to parsing, towards a totally structurally motivated approach, and the doctrine that the true structure of language is its logical structure.

And so we come to the AI approach, which represents a return to definite tasks and procedures and now slowly but surely a withdrawal from the structurally-motivated approach. Most early work on language within AI was highly structurally-motivated: the importance of the Predicate Calculus as the basic form of representation of language was taken for granted by many leading AI workers. What they were interested in was proving theorems via the derivation of one predicate calculus representation from another, an as far as they were concerned there

<u>simply was no "problem of language"</u> for the three "intractables" of language ambiguity listed above were just ignored. In the terms I developed earlier, they really did believe that natural language was a <u>projection</u> of some underlying logical format. Slowly but surely, reality has crept in. Most recently, in a survey of AI language systems, Terry Winograd (Winograd '73) has distinguished between "first" and "second" generation AI language systems, modestly placing his own at the end of the first generation. I shall not examine his distinction in any detail, but it does, I believe draw a line almost exactly where I described the emergence of AI language systems from a wholly structurally-motivated approach.

There is a long way to go yet. Some of those whose work I shall describe and whose systems would be called second generation, are still highly structuralist in approach: in that the inferential relations of the structures themselves are the centre of interest. What I shall try to maintain as we go along is that reasonings about the world, say, important as they are for understanding language content (and that this is now clear is rightly seen as a dearly won possession of the AI approach to language, the linguists having almost altogether ignored it), nonetheless there remains an important level of <u>linguistic inference</u>, not reducible to reasonings about the world. Let me give an example here to try and pin this rather abstract claim down.

We shall examine a number of examples in this course, as will others in their courses, where a pronoun in a text cannot be resolved (that is, assigned to its right referrent), unless we grant some capacity to reason about states of affairs, and general truths about people and the world. So, for example, in "The boy next door is

looking for a friend and I hope one turns up", we know immediatly that the "one" refers to a friend and not to the boy next door, and on such a correct assignment the correctness of the whole translation might rest. I will not go into details about the rules we might suggest for its resolution, but it is clearly not a merely syntactic matter as it would be if the rule were simply "assign a pronoun to the last preceding noun". I would suggest intuitively that we know the answer because of knowledge that a state of affairs wanted by one person can be expected to be one wanted, or not wanted, by another person, and that a "looking for" is expected to be succeeded by a state of "finding" or some other manifestation of what is sought. This all sounds a little abstract, but it could plausibly be put in rule form, and then a sequence of reasonings exhibited that would lead us correctly to identify "one" with the friend, rather than with the boy. You will see a number of detailed demonstrations of this sort of thing in this and other courses.

But, and this is the point, suppose the sentence had been "The boy next door is looking for a friend and I hope he finds one", then I would say that these reasonings might have been a excessive and that an intelligent program should get the right answer by simply inferring it from a repeated pattern equivalent to "PERSON SEEK PERSON", by identifying the agents and objects of the two occurrences of the pattern, thus yielding the right result. This procedure would of course be inferential, but it is not in any simple sense, reasoning about the world, or at least not directly, in the way the earlier reasoning was. ONe of the great problems, as I see it,

is how we marry these two modes together in a system that knows when to apply the right one. I shall return to this point in greater detail when discussing my own system.

The ~~important~~ thing to bear in mind in what follows is that "parsing" is being used in a non-conventional waz: it is not being used in its standard sense in mathematical and computational linguistics, where a structure is defined and the method of its application, and the comparison between different methods, is then discussed independently. That is what I referred to earlier as the "front-end" view of parsing, in which it is defined narrowly and has only marginal interest.

In these ~~notes~~ "parsing" is being used in a field in which all proposed structures are up for question: hence, they cannot be assumed but must be described in some detail first and in their own terms, as well as in terms of their adequacy to tackle what I called the funcamental problems of the systematic ambiguity of natural language. Throughout, I shall take the view, set out above, that the possibility of parsing is what gives significance to a proposed structure (rather than being merely an empirical test of it), in that a good STRUCTURE SHOWS HOW IT IS TO BE APPLIED -- not down to the very bottom level algorithms, but in general form. In these notes, the refore, you will find a lot of description of the structures offered by researchers, and of comparisons between them, and not too much of the actual nitty-gritty of the sequential forms through which the representation of a sentence passes in the actual parsing process. The reasons for this are simple: in the case of some of

the best known representation systems for English, there is no effective parsing system as yet, and so to ignore those would truncate the coverage of the course in an arbitrary way. Again, for others, such as Schank's, there is a parser in operation, but since Schank's system is always presented in a wholly static non-operational way, discussions of the structure and its parser are wholly separate, and it is clear which one must be described in notes when a competition for space arises as it does here.

However, there will be detailed description of actual, blow-by-blow, processing of sentences, based on the descriptions and comparisons set out here. These will be done by means of transparancies, and will include at least a display of Winograd's parsings (concentrating particularly on their semantic aspects), and of Riesbeck's parser for Schank's structures, and of my own parser. I will also try and introduce some details of the actual size of the implementations in terms of vocabulary and classes of sentences actually processed: a matter it is not always easy to elicit from the published literature. All I can promise is that I will try to look with no more jaundiced an eye at the systems of others than I do at my own.

2. WINOGRAD'S UNDERSTANDING SYSTEM

Winograd's system runs as a dialogue, in real time, between a human operator who types in messages, and the system proper, called SHRDLU, which displays on a screen pictures of a closed world of coloured blocks and pyramids, a box into which they can be put, and an arm or crane that

can move the objects about. SHRDLU displays its replies to the operator by writing on the same screen. The following is a discontinuous sample of the dialogue Winograd gives in his book. SHRKLU'S replies are in upper case:

pick up the big red block

OK

find a block which is taller than the one you are holding and put it into the box.

BY 'IT', I ASSUME YOU MEAN THE BLOCK WHICH IS TALLER THAN THE ONE I AM HOLDING

OK

what does the box contain

THE BLUE PYRAMID AND THE BLUE BLOCK

had you touched any pyramid before you put the green one on the little cube

YES' THE GREEN ONE

when did you pick it up

WHILE I WAS STACKING UP THE RED CUBE' A LARGE RED BLOCK AND A LARGE GREEN CUBE.

why

TO GET RID OF IT.

This ingenious program was also able to accept new definitions of words, such as 'steeple'; to make use of what is knew about the properties of the objects described, such as that a block could have its top cleared off, but a pyramid could not; and also to remember what it had done before, as in the sample above.

The program was written in the language PLANNER. which is a concrete expression of the slogan "meanings are procedures", a sentiment into whose own meaning it is probably best not to inquire too closely, but which has

undoubtedly led to a new style of programming. PLANNER[$] is a theorem proving language: it tries to establish the truth of assertions, not in the normal uniform, proof-theoretic manner, but by accepting a range of "programmed hints" about how best to proceed at any point. In a language understanding program like Winograd's, this means replacing familiar grammar rules such as S → NP + VP. (a sentence consists of a noun phrase followed by a verb phrase) by procedures, in this case:

```
((PDEFINE SENTENCE ((PARSE NP) NIL FAIL)((PARSE VP)FAIL
FAIL RETURN)))
```

The details of the notation need not detain us at this point; what is important is that Winograd's grammar is not the conventional list of rules, but small subprograms like the lines above, that actually represent <u>procedures for imposing</u> the desired grammatical structure.

The syntactic structure to be imposed is Halliday's (1970) hierarchical system of clauses, that can be illustrated as follows:

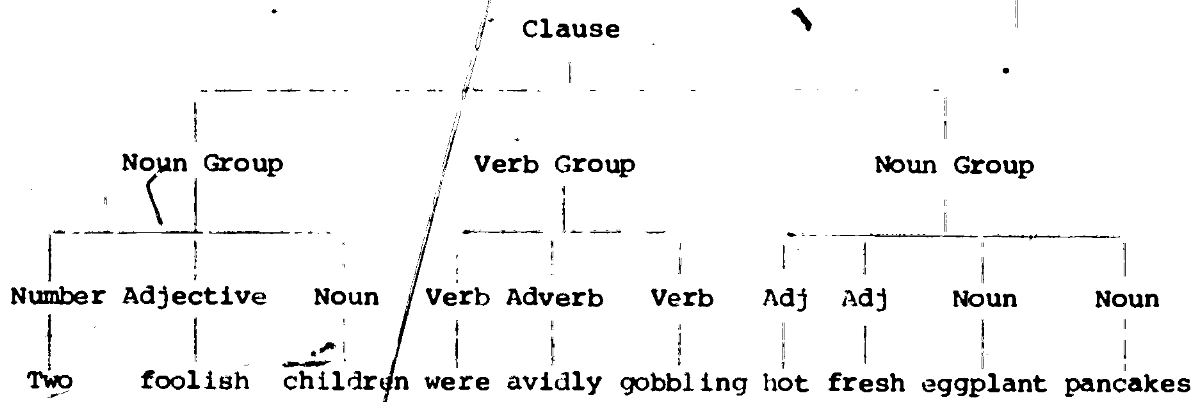

$ Actually the grammar was written in a language PROGRAMMAR, and the semantic and inference parts of the program in LISP and a subset of PLANNER called MICRO-PLANNER.

In terms of the notions set out earlier Winograd's parsing is top down, and depth first, with no automatic back up. The parsing program for each grammatical category is a functional definition in PROGRAMMAR, which can be stated either as above for SENTENCE, or as a flow-chart as below for VP:

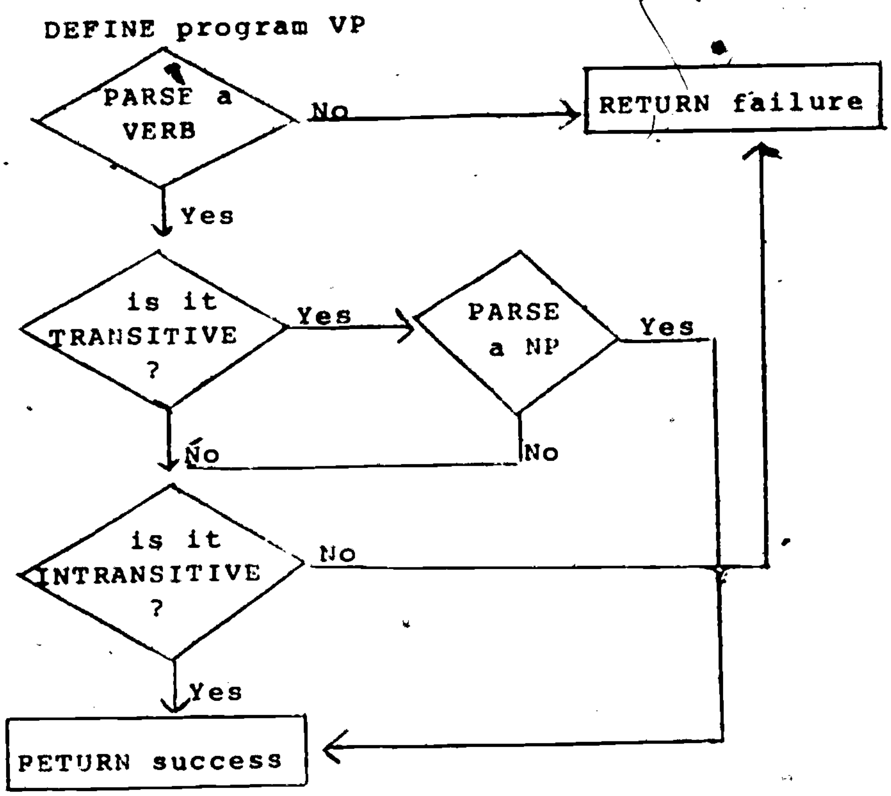

Here is Winograd's own account of the start of this top-down parsing procedure for the sentence "Pick up a big red block":

"The CLAUSE program looks at the first word, to decide what unit the CLAUSE begins with. If it sees an adverb, it assumes the sentence begins with a single-word modifier; if it sees a preposition, it looks for an initial PREPG. If it sees a BINDER, it calls the CLAUSE program to look for a BOUND CLAUSE. In English (and possibly all languages) the first word of a construction often gives a very good

clue as to what that construction bill be. In this case, "pick" is a verb, and indicates that we may have an IMPERATIVE CLAUSE. The program starts the VG program with the initial VG feature list (VG IMPER), looking for a VG of this type. This must either begin with some form of the verb "do", or with the main verb itself. Since the next word is not "do" it checks the next word in the input (in this case still the first word) to see whether it is the infinitive form of a verb. If so, it is to be attached to the parsing tree, and given the additional feature MVB (main verb). The current structure can be diagrammed as:

```
(CLAUSE MAJOR)
    (VG IMPER)
        (VB MVB INF TRANS VPRT---------pick
```

TRANS AND VPRT came from the definition of the word "pick" when we called the function PARSE for a word."

After the syntactic parsing, a number of "semantic specialists" attach semantic structures to specific syntactic ones. A semantic definition of an action like "grasp" would be of the form

```
(CMEANS((((#ANIMATE))((#MANIP)))
(#EVAZ (COND((PROGRESSIVE)
(QUOTE(#GRASPING #2 *TIME)))
(T(QUOTE(#GRASP #2 *TIME
        )))))NIL))
```

which says essentially that grasping is something done by an animate entity to a manipulable one (first line). More of the real content of such actions is found in their inferential definition. Here is the one for "pickup":

```
(DEFTHEOREM TC-PICKUP(THCONSE (X(WHY (EV))EV)
    ("PICKUP $?X) (MEMORY) (THGOAL(#GRASP $?X) (THUSE TC-GRASP))
    (THGOAL (#RAISEHAND) (THNODB) (THUSE TC-RAISEHAND))
    (MEMOREND(#PICKUP $?EV  $?X))))
```

Once again the details of the notation nedd not be explained in order to see that the word is being defined in terms of a number of more primitive sub-actions, such as RAISEHAND' each of which must be carried out in order that something may indeed be picked up.

In the case of "a red cube", the following structure is built up by an NP "semantic specialist"

```
        (GOAL (IS X BLOCK))
        (GOAL (COLOR X RED))
        (EQDIM X---------------------PLANNER description
    (BLOCK MANIP PHYSOB THING)----------- markers
```

The first three lines are procedures that when evaluated will seek an object X that is a block, is equidimensional

(EQDim) and is red. The last line is a set of "semantic features" read off right to left from the following "feature tree"

```
                |NAME
                |PLACE          |SHAPE
                |PROPERTY---|SIZE
                |               |LOCATION
                |               |COLOR          |ROBOT
                |ANIMATE----------------------|
                |                               |HUMAN
                |                   |BLUE
                |                   |RED
                |          (------ |BLACK
                |          (       |WHITE
                |          (       |GREEN                              |STACK
   THING---     |PHYSOB--  (   |CONSTRUCT------------------------|PILE
                |          (   |HAND                                 |ROW
                |          (-- |TABLE          |PYRAMID
                |              |MANIP---------|BLOCK
                |                              |BALL
                |
                |RELATION-------------|EVENT
                                      |
                                      |TIMELESS
```

This whole semantic structure can be used by the deductive component at the system, before evaluation resulting in the actual picking up, to see if such an object is possible. If it were not, (an "equidimensional pyramid" would not be), the system could go back and try to re-parse the sentence.

One reason for the enormous impact of this work was that, prior to its appearance, AI work was linguistically trivial, while the systems of the linguists had no place for the use of inference and real world knowledge. Thus a very limited union between the two techniques was able to breed considerable results. Before Winograd there were few programs in AI that could take a reasonable complex English sentence and ascribe any structure whatever to it. In early classics of 'natural language understanding' in AI, such as Bobrow's

STUDENT (1968) problem-solver for simple algebra, input sentences had to be short and of stereotyped form, such as "what is the sum of.......?"

Conversely, in linguistics, there was, until very recently, little speculation on how we understand the reference of promouns in such elementary sentences "the soldiers fired at the women and I saw several fall", where it is clear that the answer is both definite, and that finding it requires some inferential manipulation of generalizations about the world. The reader should ask himself at this point how he knows the referent of the pronoun in that sentence.

3. SOME DISCUSSION OF SHRDLU

So far, the reaction to Winograd's work has been wholly uncritical. What would critics find to attack if they were so minded? Firstly, that Winograd's linguistic system is highly conservative, and that the distinction between 'syntax and 'semantics' may not be necessary at all. Secondly, that his semantics is tied to the simple referential world of the blocks in a way that would make it inextensible to any general, real world, situation. Suppose 'block' were allowed to mean 'an obstruction' and 'a mental inhibition', as well as 'a cubic object'. It is doubtful whether Winograd's features and rules could express the ambiguity, and, more importantly, whether the simple structures he manipulated could decide correctly between the alternative meanings in any given context of use. Again, far more sophisticated and systematic case structures than those he used might be needed to resolve the ambiguity of 'in' in 'He ran the mile in five minutes', and 'He ran the mile in a paper bag', as well as the combination of case with word sense ambiguity

in 'He put the key in the lock' (door lock) and 'He threw the key in the lock" (river lock).

The blocks world is also strongly deductive and logically closed. If gravity were introduced into it, then anything supported that was pushed in a certain way would have, logically have, to fall. But the common sense world, of ordinary language, is not like that: in the 'women and soldiers' example given earlier, the pronoun 'several' can be said to be resolved using some generalisation such as 'things shot at and hurt tend to fall'. There are no logical 'have to's' there, even though the meaning of the pronoun is perfectly definite.

Indeed, it might be argued that, in a sense, and as regards its semantics, Winograd's system is not about natural language at all, but about the other technical question of how goals and sub-goals are to be organised in a problem-solving system capable of manipulating simple physical objects. If one glances back at the definition of 'pickup' quoted above, one can see that it is in fact an expression of a procedure for picking up an object in the SHRDLU system. Nothing about it, for example, would help one understand the perfectly ordinary sentence 'I picked up my bags from the platform and ran for the train', let alone any sentence not about a physical action performable by the hearer. One could put the point so: what we are given in the PLANNER code is not a sense of 'pick up' but a case of its use, just as 'John picked up the volunteer from the audience by leaning over the edge of the stage and drawing her up by means of a rope clenched in his teeth' is not so much a sense of the verb as a use of it.

Those who like very general analogies may have noticed that Wittgenstein (1953, para. 2ff.) devoted considerable space to the construction of an elementary language of blocks, beams and slabs; one postulated on the assumption that the words of language were basically, as is supposed in model theory, the names of items. But, as he showed of the enterprise, and to the satisfaction of many readers, "That philosophical concerpt of meaning (i.e. of words as the unambiguous names of physical objects---YW) has its place in a primitive idea of the way language functions. But one can also say that it is the idea of a language more primitive than ours". (my italics). Thus in terms of the notions of the introduction one might say that there is a strong structurally (rather than content) motivated element in SHRDLU.

In this course I do not discuss the excellent and widely known work of Woods. Let me explain why. The system, based on an augmented state transition network grammar, is undoubtedly one of the most robust in actual use, in that it is less sensitive to the PARTICULAR input questions it encounters than its rivals. The reason for not treating it in depth is that both Woods and Winograd have argued in print that their two systems are essentially equivalent (Winograd 1971) (Woods 1973), and so, if they are right, there is no need to discuss both, and Winograd's is, within the AI community at least, the better known of the two.

Their equivalence arguments are probably correct: both are grammar-based deductive systems, operating within a question-answering environment in a highly limited domain of discourse. Winograd's system of hints on how to proceed,

within his PROGRAMMAR grammar, is, as he himself point out, formally equivalent to an augmented state transition network, and in particular to the ordering of choices at nodes in Woods' system.

There is a significant difference in their metaphysical approaches, or presuppositions about meaning which, however, has no influence on the actual operation of their respective systems. This difference is disguised by the allegiance both give to a 'procedural view of meaning'. The difference is that Woods takes a much more logico-semantic interpretation of that slogan than does Winograd. In particular, for Woods the meaning of an input utterance to his system is the procedures within the system that manipulate the truth conditions of the utterance and establish its truth value. To put the matter crudely, for Woods' an assertion has no meaning if his system cannot establish its truth or falsity. Winograd has certainly not committed himself to any such extreme position.

It is interesting to notice that Woods' is, in virtue of his strong position on truth conditions, probably the only piece of work in the field of AI and natural language to satisfy Hayes' (1974) recent demand that to be "intellectually respectable" a knowledge system must have natural model theoretic semantics, in Tarski's sense. Since no one has ever given precise truth conditions for any interesting piece of discourse, such as, say, Woods' own papers, one might claim that his theoretical presuppositions necessarily limit his work to the analysis of micro-worlds (as distinct from everyday language).

There is a low-level problem about the equivalance of Woods' and Winograd's systems, if we consider what we might call the received common-sense view of their work. Consider

the following three assertions:

(1) Woods' system is an implementation of a transformational grammar.

(2) Winograd's work has shown the irrelevance of transformational grammar for language analysis--a view widely held by reviews of his work.

(3) Woods' and Winograd's systems are formally equivalent--a view held by both of them.

There is clearly something of an inconsistent triad amongst those three widely held beliefs. The trouble probably centers on the exact sense which Woods' work is formally equivalent to a transformational grammar--not a question that need detain us here, but one worth pointing out in passing.

4. SECOND GENERATION SYSTEMS

To understand what was meant when Winograd contrasted his own with what he called second generation systems, we have to remember, as always in this subject, that the generations are of fashion, not chronology or inheritance of ideas. He described the work of Simmons, Schank and myself among others in his survey of new approaches, even though the foundations and terminology of those approaches were set out in print in 1966, 1967 and 1968 respectively. What those approaches, and others, have in common is the belief that understanding systems must be able to manipulate very complex linguistic objects, or semantic structures, and that no simplistic approaches to understanding language with computers will work.

In an unpublished, but already very influential recent paper, Minsky (1974) has drawn together strands in the work of Charniak (1972) and the authors above using a terminology of 'frames':

"A frame is a data-structure for representing a stereotype situation, like a certain kind of living room, or going to a children's birthday party. Attached to each frame are several kinds of information. Some of this is information about how to use the frame. Some is about what one can expect to happen next. Some is about what to do if these expectations are not confirmed.

We can think of a frame as a network of nodes and relations. The top levels of a frame are fixed and represent things that are always true about the supposed situation. The lower levels have many terminals---'slots' that must be filled by specific instances or date. Each terminal can specify conditions its assignments must meet.... Simple conditions are specified by markers that might require a terminal assignment to be a person, an object of sufficient value, etc...."

The key point about such structures is that they attempt to specify in advance what is going to be said, and how the world encountered is going to be structured. The structures, and the inference rules that apply to them, are also expressions of ' partial information' (in McCarthy's phrase) that are not present in first generation systems. As I showed earlier, with the 'women and soldiers' example, such loose inductive information, seeking confirmation from the surrounding context, is required for very simple sentences. In psychological and visual terms, frame approaches envisage an understander as at least as much a looker as a seer.

I shall now describe briefly three approaches that might be called second generation; Colby, Schank and myself. I shall not discuss Charniak's work here, largely because he is describing it himself, in another course, but partly because he does not implement his system with a parser. He also disagrees with the view of parsing I sketched in Section I -- namely that a representational structure gets much of its significance from its actual or potential parsing, and that parsing therefore cannot be decoupled--I shall return to his work in the later comparisons and assume some knowledge of it by my audience.

Colby

Colby has (Colby et al '71), (Enea and Colby '73), (Colby and Parkison, '74) developed over a period of several years an on-line dialogue system that simulates a paranoid patient called PARRY. A measure of its success is that it is perhaps the most used AI program in existence; used, that is, by those outside the research project that developed it. Here is a discontinuous sample of its output, again with the system's responses in the upper case:

Have you been hospitalised before?

THIS IS THE FIRST TIME.

How long have zou been there?

ABOUT TWO WEEKS.

Any headaches?

MY HEALTH IS FINE.

Are you having any difficulty with your memory?

NO.

Then perhaps you can tell me who is the President of the United States?

NIXON IS PRESIDENT.

What day is it today?

TODAY IS THURSDAY.

What season is it?

LETS GET ON WITH THE INTERVIEW.

Colby's aid was to produce a system that was not, in his word, 'fragile' in the way that most AI systems are: one that did not break down, or give up, if the input was not exactly what it expected; but one that carried on somehow in all circumstances, in the way that conversing humans do.

Fragility is avoided by having no conventional syntax analyser, and by not even attempting to take account of all the words in the input. This is a considerable aid, since any parser that begins to parse a more than usually polite request such as 'Would you be so kind as to...' is going to be in trouble. British English speakers arriving in the U.S. quickly learn to delete such phrases, since they cause great confusion to human listeners in stores.

The input text is segmented by a heuristic that breaks it at any occurrence of a range of key words. Patterns are then matched with each segment. There are at present about 1700 patterns on a list (Colby and Parkison, in press) that is stored and matched, not against any syntactic or semantic representations of words (except to deal with contractions and misspellings), but against the input word string direct, and by a process of sequential deletion. So, for example, 'What is your main problem' has a root verb 'BE' substituted to become

WHAT BE YOU MAIN PROBLEM.

It is then matched successively in the following forms after successive deletions:

BE YOU MAIN PROBLEM
WHAT YOU MAIN PROBLEM
WHAT BE MAIN PROBLEM
WHAT BE YOU PROBLEM
WHAT BE YOU MAIN

and only the penultimate line exists as one of the stored patterns and so is matched. Stored in the same format as the patterns are rules expressing the consequences for the 'patient' of detecting aggression and over-friendliness in the interviewer's questions and remarks. The matched patterns round are then tied directly, or via these inference rules, to response patterns which are generated.

Enormous ingenuity has gone into the heuristics of this system, as its popularity testifies. The system has also changed considerably: it is now called PARRY 2 and contains the above pattern matching, rather than earlier key work, heuristics. It has the partial, or what some would call 'pragmatic' , rules about expectation and intention, and these alone might qualify it as 'second generation' on some interpretations of the phrase. A generator is also being installed to avoid the production of only 'canned' responses.

Colby and his associates have put considerable energy into actually trying to find out whether or not psychiatrists can distinguish PARRY'S responses from those of a patient (Colby and Hilf '73). This is probably the first attempt to actually apply Turing's test of machine-person distinguishability. There are statistical difficulties about interpreting the results but, by and large, the result is that the sample questioned cannot distinguish the two. Whether or not this will influence those who still, on principle, believe that PARRY is not a simulation because it

'does not understand', remains to be seen. It might be argued that they are in danger of falling into a form of Papert's 'human-superhuman fallacy': of attacking machine simulations because they do not perform superhuman tasks, like translate poetry, tasks that some people certainly can do. When such sceptics say that PARRY does not understand they have in their minds a level of understanding that is certainly high - one could extend their case ironically by pointing out that very few people understand the content of sentences in the depth and detail that an analytic philosopher does, and a very good thing too. There can be no doubt that many people on many occasions DO seem to understand in the way that PARRY does.

Schank

Schank's is a rich system of semantic representation, developed over a period of six years, with the collaboration of a number of talented students. Its graph system of notation has influenced psychologists like Abelson (1973), among others. Schank's contribution has been the notational system representing the structure of natural language sentences, and this has been programmed by various collaborators over the years. In its present version, called MARGIE (Schank et al '73) it has an analyser of English due to Riesbeck (1974), a semantic memory component due to Rieger (1974), and a generator of English due to Goldman (1974). MARGIE produces output in two modes, demonstrating the sort of conceptual inferencing that goes on at the level of the semantic$ representation: the PARAPHRASE mode and the INFERENCE mode. Samples of input and output to and from the two modes can be shown thus:

$ (Footnote on page 31)

(INFERENCE MODE)

INPUT: John told Mary that Bill wants a book.

OUTPUT: A book about what

Mary knows that Bill wants a book.

Bill wants to come to have a book.

Bill wants someone to cease to have a book.

(PARAPHRASE MODE)

INPUT: John killed Mary by choking Mary.

OUTPUT: John strangled Mary.

John choked Mary and she died because she could not breathe.

Mary died because she was unable to inhale some air and she was unable to inhale some air because John grabbed her neck.

The aim of Schank's system has always been to provide a representation of meaning in terms of which these and other tasks, such as machine translation, could be carried out; a representation, moreover, that was to be independent of any particular languages, and of syntax, and indeed, of all surface structure whatever.

The formal structure of Schank's graphs is that of dependency grammar (Hays '64), and the items in the graph are of four types, or conceptual categories. They are symbolised as PP, ACT, PA and AA, which are acronyms, but which correspond closely (for the purpose of understanding their

---

$ Schank distinguishes 'conceptual' and 'semantic' representations in a way that is important for him within his own system. However, I shall use the terms indifferently since, in this brief and superficial description, nothing hangs upon the distinction.

function) to those of a noun, verb, adjective and adverb, respectively. The basic structure is called a conceptualisation, and is normally introduced with a straightforward dependency structure such as, for the sentence 'The man took a book':

$$\text{Man} \overset{p}{\Longleftrightarrow} \text{take} \xleftarrow{o} \text{book}$$

Here 'p' indicates past, and is the dependency symbol linking a PP to the ACT ('take') which is the hub of the conceptualisation, as with Simmons.* The 'o' indicates the objective case, marking the dependence of the object PP on the central ACT. There is a carefully constructed syntax of linkages between the conceptual categories, that will be described only in part in what follows:

The next stage of the notation involves an extended case notation and a set of primitive ACTs, as well as a number of items such as PHYSCONT which indicate other states, and items of a fairly simplified psychological theory (the dictionary entry for advise, for example, contains a subgraph telling us that Y ' will benefit' as part of the meaning of 'X advises Y' (Schank '73)). There are four cases in the system, and their subgraphs are as follows:

Objective case:	$\text{ACT} \xleftarrow{o} \text{PP}$
Recipient case+	$\text{ACT} \xleftarrow{R} \begin{cases} \rightarrow \text{PP} \\ \leftarrow \text{PP} \end{cases}$

* This is a considerable oversimplification, made in order to give a brief and self contained description. But, in fact, many English nouns are represented as ACT's: chair, pen, honesty, and transportation.

Instrumental case:

Directive case:

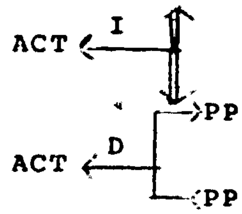

There are at present fourteen* basic actions forming the nubs of the graphs, as well as a default action DO. They are: PROPEL, MOVE, INGEST, EXPEL, GRASP, PTRANS, MTRANS, ATRANS, SMELL, SPEAK, LOOK-AT, LISTEN-TO, CONC and MBUILD. The notions of case and primitive act are related by rules in the development of conceptualisations. So, for example, the primitive act INGEST has as its instrument the act PTRANS. There are also other inferences from any ACT classified as an INGEST action, such as that the thing ingested changes its form; that if the thing ingested is edible the ingester becomes ' more nourished' etc. (see Schank '73, pp. 38ff.). This will all become clearer if we consider the transition from a dictionary entry for an action to a filled-in conceptualisation. Here is the dictionary entry for the action 'shoot':

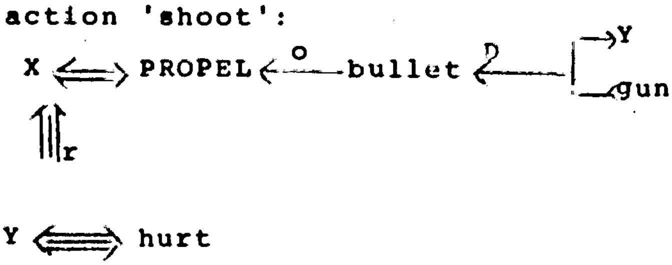

We can consider this entry as an active 'frame-like' object seeking filler items in any context in which it is activated. Thus, is the sentence 'John shot the girl with a rifle', the variables will be filled in from context and the case inference will be made from the main act PROPEL, which is that its instrument is MOVE' GRASP or PROPEL, and so we will arrive at the whole conceptualisation:

* Since the publication of (Schank 73a) their number has been reduced to eleven (plus DO) by the elimination of SMELL' LISTENTO' LOOKAT and CONC, and the addition of ATTEND.

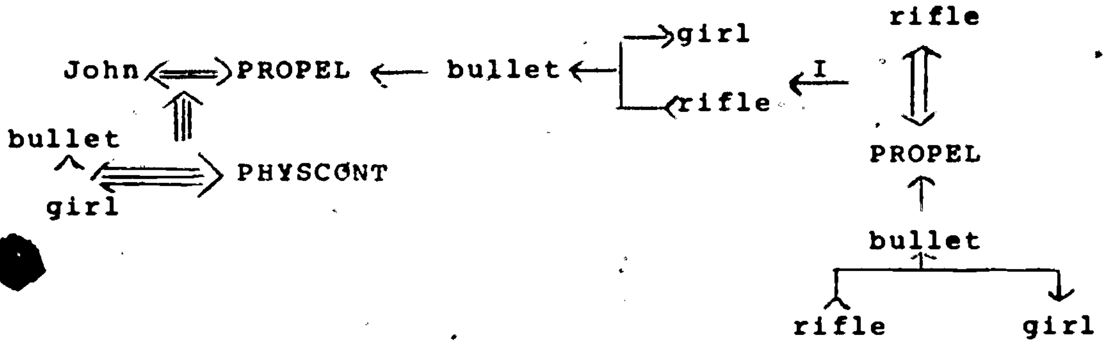

This case inference must be made, according to Schank, in order to achieve an adequate representation. There is, in the last diagram, a certain redundancy of expression, but as we shall see in the next section this often happens with deeper semantic notations.

More recently, Schank, together with Rieger, has developed a new class of causal inferences which deepen the diagrams still further. So, in the analysis of 'John's cold improved because I gave him an apple' (in Schank '74a) the extended diagram contains at least four yet lower levels of causal arrowing, including one corresponding to the notion of John constructing the idea (MBUILD) that he wants to eat an apple. So we can see that the underlying explication of meaning here is not only in the sense of linguistic primitives, but in terms of a theory of mental acts as well.

NOw there are a number of genuine expositional difficulties here for the commentator faced with a system of this complexity. One aspect of this is the stages of development of the system itself, which can be seen as a consistent process of producing what was argued for in advance. For example, Schank claimed early on to be a constructing system of semantic structures underlying the 'surface of natural language', although initially there were no primitives at all, and as late as (Schank et al '70) there was only a single primitive TRANS, and most of the entries in

the dictionary consisted of the English words coded, together with subscripts. Since then the primitive system has blossomed and there are now twelve primitives for ACTS including three for the original TRANS itself. Each exposition of the system recounts its preceding phases, from the original primitive-free one, through to the present causal inference form; rather as each human foetus is said to relive in the womb all the evolutionary stages of the human race. The only trouble with this, from an outsider's point of view, is that at each stage the representation has been claimed, in firm tones, to be the correct one, while at the same time Schank admits, in moments of candor (Schank '73), that there is no end to the conceptual diagramming of a sentence. This difficulty may well reflect genuine problems in language itself, and, in its acutest form concerns a three-way confusion between an attractive notation for displaying the 'meanings of words', the course of events in the real world, and, finally, actual procedures for analysis. It is not always clear whether or not proocedures implementing conceptual dependenca are intended to recapture all the many phases of expansion of the diagrams.

This raises the, to me, important question of the application of a semantic system, that I shall touch on again later. Schank, for example, does mention in passing the questions of word-sense ambiguity, and the awful ambiguity of English prepositions, but they are in no way central for him, and he assumes that with the availability of 'the correct representation', his system when implemented must inevitably solve these traditional and vexing questions. No procedures are hinted at along with the graphs as to how this is to be done. A distinction of importance may be becoming

apparent here between Schank's work and Rieger's: in Rieger's thesis (Rieger '74) the rules of inference appear to create separate and new subgraphs which may stand in an inferential relation to each other so as to produce conclusions about problems of, say, pronoun reference, etc. But in Schank's corresponding papers the same inferences are not applied to actual problems (Schank '74a) but only complexify the conceptual graphs further. Closely connected with this is the question of the survival of the surface structure in the diagrams. Until very recently primitiveisation applied only to verbs, that of nouns being left to Weber (Weber '72). Most recently, though, noun words have been disappearing from diagrams and been replaced by categories such as 'PHYSOBS'. But it is clear that the surface is only slowly disappearing, rather than having been abhorred all along.

In his most recent publication (Schank '74b) there are signs that this trend of infinitely proliferating diagrams is reversing. In it Schank considers the application of his approach to the representation of text, and concludes, correctly in mn view, that the representations of parts of the text must be interconnected by causal arrows, and that, in order to preserve lucidity, the conceptual diagrams for individual sentences and their parts must be abbreviated, as by triples such as PEOPLE PTRANS PEOPLE. Here indeed, the surface simply has to survive in the representation unless one is prepared to commit oneself to the extreme view that the ordering of sentences in a text is a purely superficial and arbitrary matter. The sense in which this is a welcome reversal of a trend should be clear, because in the 'causation-inference' development, mentioned earlier, all the consequences and effects of a conceptualization had to be drawn.

within itself. Thus, in the extreme case, each sentence of a text should have been represented by a diagram containing most or all of the text of which it was a part. Thus the representation of a text would have been impossible on such principles.

At this point I shall describe in some detail with slides the implementation of Schank's system as a parser by Riesbeck.

Wilks

My own system also has a uniform representation, in terms of structures of primitives, for the content of natural language. It is uniform in that information that might conventionally be considered syntactic, semantic or factual is all represented within a single structure of complex entities* all of which are in turn constructed from a budget of 80 primitive semantic entities.

The system runs on-line as a package of LISP and MLISP programs, taking as input small paragraphs of English, that can be made up by the user from a vocabulary of about 600 words, and producing a good French translation as output. This environment provides a pretty clear test of language understanding, because French translations for everyday prose are either right or wrong, and can be seen to be so, while at the same time, the major difficulties of understanding programs ------- word sense ambiguity, case ambiguity, difficult pronoun reference, etc. ------- can all be represented within a machine translation environment by, for example, choosing the words of the input sentence containing a pronoun reference difficulty so that the possible alternative references have different genders in French. In that way the French output makes quite clear whether or not the

*called formulas and paraplates

program has made the correct inferences in order to understand what it is translating. The program is reasonably robust in actual performance, and will even tolerate a certain amount of bad grammar in the input, since it is not performing a syntax analysis in the conventional sense, but seeking messages representable in the semantic structures employed.

Typical input would be a sentence such as "John lives out of town and drinks his wine out of a bottle. He then throws the bottles out of the window." The program will produce French sentences with different output for each of the three occurrences of "out of", since it realises that they function quite differently on the three occasions of use, and that the difference must be reflected in the French. A sentence such as "Give the monkeys bananas although they are not ripe because they are very hungry" produces a translation with different equivalents for the two occurrences of "they", because the system correctly realises, from what I shall describe below as <u>preference</u> considerations, that the most sensible interpretation is one in which the first "they" refers to the bananas and the second to the monkeys, and bananas and monkeys have different genders in French. These two examples are dealt with in the "basic Mode" of the system (Wilks '73a). In many cases it cannot resolve pronoun ambiguities by the sort of straightforward "preference considerations" used in the last example, where, roughly speaking, "ripeness" prefers to be predicated of plant-like things, and hunger of animate things. Even in a sentence as simple as "John drank the wine on the table and it was good", such considerations are inadequate to resolve the ambiguity of "it" between wine and table, since both may be good things. In such cases, of inability to resolve within

its basic mode, the program deepens the representation of the text so as to try and set up chains of inference that will reach, and so prefer, only one of the possible referents. I will return to these processes in a moment, but first I shall give some brief description of the basic representation set up for English.

For each sense of a word in its dictionary the program sees a formula. This is a tree structure of semantic primitives, and is to be interpreted formally using dependency relations. The main element in any formula is the rightmost, called its head, and that is the fundamental category to which the formula belongs. In the formulas for actions, for example, the head will always be one of the primitives PICK, CAUSE, CHANGE, FEEL, HAVE, PLEASE, PAIR, SENSE, USE, WANT, TELL, BE, DO, FORCE, MOVE, WRAP, THINK, FLOW, MAKE, DROP, STRIK, FUNC or HAPN.

Here is the tree structure for the action of drinking+

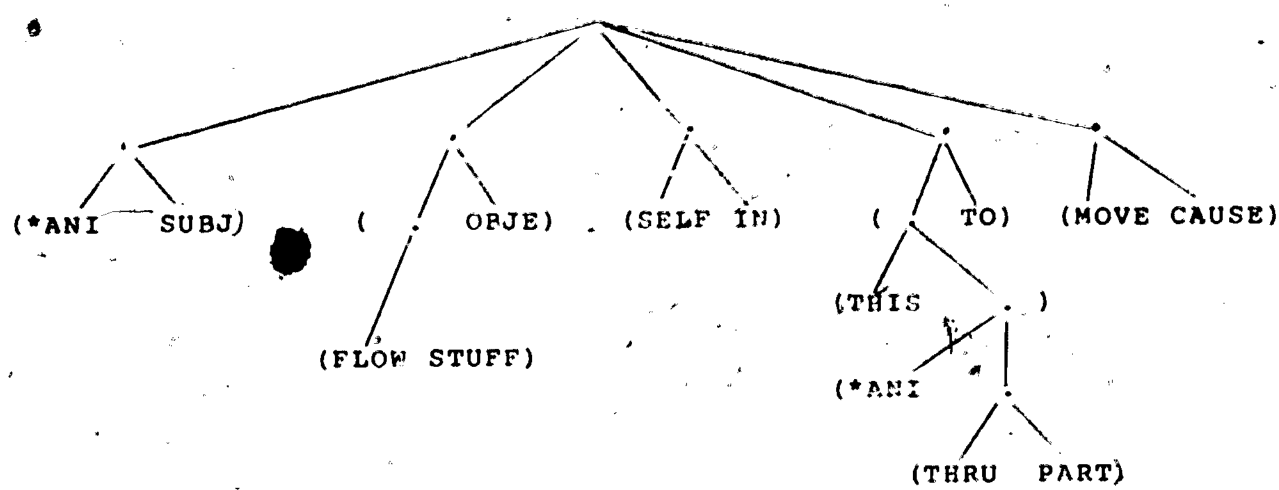

Once again, it is not necessary to explain the formalism in any detail, to see that this sense of "drink" is being expressed as a causing to move a liquid object (FLOW STUFF) by an animate agent, into that same agent (containment case

indicated by IN, and formula syntax inditifies SELF with the agent) and via (direction case) an aperture (THRU PART) of the agent.

Template structures, which actually represent sentences and their parts are built up as networks of formulas like the one above.* Templates always consist of an agent node, and action node and an object node, and other nodes that may depend on these. So, in building a template for "John drinks wine", the whole of the above tree-formula for "drinks" would be placed at the action node, another tree structure for 'John' at the agent node and so on. The complexity of the system comes from the way in which the formulas, considered as active entities, dictate how other places in the same template should be filled.

Thus, the "drink" formula above can be thought of an an entity at a template action node, seeking a liquid object, that is to saw a formula with (FLOW STUFF) as its right-most branch, to put at the object node of the same template. This seeking is preferential, in that formulas not satisfying that requirement will be accepted, but only if nothing satisfactory can be found. The templage <u>finally established</u> for a fragment of text is the one in which the most formulas have their preferences satisfied. There is a general principle at work here: that the right interpretation "says the least" in information-carrying terms. This very simple device is able to do much of the work of a syntax and word-sense ambiguity resolving program. For example, if the sentence had been "John drank a whole Pitcher", the formula for the "pitcher of liquid" would have been preferred to that for the human, since the subformula (FLOW STUFF) could be appropriately located

within it.

A considerable amount of squeezing of this simple canonical form of template is necessary to make it fit the complexity of language: texts have to be fragmented initially; then, in fragments which are, say, prepositional phrases there is a dummy agent imposed, and the prepositional formula functions as a pseudo-action. There are special "less preferred" orders to deal with fragments not in agent-action-object order, and so on.

When the local inferences have been done that set up the agent-action-object templates for fragments of input text, the system attempts to tie these templates together so as to provide an overall initial structure for the input. One form of this is the anaphora tie, of the sort discussed for the "monkeys and bananas" example above, but the more general form is the case tie. Assignment of these would result in the template for the last clause of "He ran the mile in a paper bag" being tied to the action node of the template for the first clause "He ran the mile", and the tie being labelled CONTainment. These case ties are made with the aid of another class of ordered structures called <u>paraplates</u>, that are attached to the formulas for English prepositions. So, for "outof" there would be at least six ordered paraplates, each of which is a string of functions that seek inside templates for information. In general, paraplates range across two, not necessarily contiguous, templates. So, in analysing "He put the number he thought of in the table", the successfully matching paraplate would pin down the dependence of the template for the last of the three clauses as DIREction, by taking as argument only that one template for the last clause that contained the formula

for a numerical table, rather than a kitchen table, in virtue, in this example, of the function in that paraplate seeking a similarity of head (SIGN in this case) between the two object formulas, for "number" and "table".

The structure of mutually connected templates that has been put together thus far constitutes a "semantic block", and if it can be constructed, then, as far as the system is concerned, all semantic and referential ambiguity has been resolved and it will begin to generate Franch by unwrapping the block again. The generation aspects of this work have been described in (Herskovits). One aspect of the general notion of preference is that the system should never construct a deeper or more elaborate semantic representation than is necessary for the task in hand and, if the initial block can be constructed and a generation of French done, no "Deepening" of the representation will be attempted.

However, many examples cannot be resolved by the methods of this "basic mode" and, in particular, if a word sense ambiguity, or pronoun reference, is still unresolved, then a unique semantic block of templates cannot be constructed and the "extended mode" will be entered.* In this mode, new template-like forms are extracted from existing ones, and then added to the template pool from which further inferences can be made. So, in the template derived earlier for "John drinks wine", the system enters the formula for "drinks", and draws inferences corresponding to each case sub-formula. In this example it will derive template-like forms equivalent to, in ordinary English, "The wine is in John", "The wine entered John via an aperture" and so on. The extracted templates express information already implicitly present in the text, even though many of them are

*Wilks '73b, and in press

partial inferences: ones that may not necessarily be true.

Common-sense inference rules are then brought down, which attempt, by a simple strategy, to construct the shortest possible chain of rule-linked template forms from one containing an ambiguous pronoun, say, to one containing one of its possible referents. Such a chain then constitutes a solution to the ambiguity problem, and the preference approach assumes that the shortest chain is always the right one. So, in the case of "John drank the wine on the table and it was good", the correct chain to "wine" uses the two rules:

I 1. ((*ANI 1) ((SELF IN) (MOVE CAUSE)) (*REAL 2))→(1(*JUDG) 2)

or, in "semi-English",

[animate 1 cause-to-move-in-self real-object-2]→[1 *judges 2]

I 2. (1 BE (GOOD KIND))←→((*ANI 2) WANT 1)

or, again,

[1 is good]←→[animate-2 wants 1]

These rules are only partial, that is, they correspond only to what we may reasonably look out for in a given situation, not to what MUST happen. The hypothesis here is that understanding can only take place on the basis of simple rules that are confirmed by the context of application. In this example the chain constructed may be expressed (using the above square bracket notation to contain not a representation, but simply an indication, in English, of the template contents) as:

forwards inf. ↓	[John drank the wine]		Template 1
	[John causes-to-move-in-self	wine]	⁝ Template 1
	[John *judges	wine]	by I1.
backwards inf. ↑	[John wants	wine]	⁝ line above
	[wine is	good]	by I2.
	[?it is	good]	Template 3

The assumption here is that no chain using other inference rules could have reached the "table" solution by using less than two rules.

The chief drawback of this system is that codings consisting entirely of primitives have a considerable amount of both vagueness and redundancy. For example, no reasonable coding in terms of structured primitives could be expected to distinguish, say, "hammer" and "mallet". That may not matter provided the codings can distinguish importantly different senses of words. Again, a template for the sentence "The shepherd tended his flock" would contain considerable repetition, each node of the template trying, as it were, to tell the whole story by itself. Whether or not such a system can remain stable with a considerable vocabulary, of say several thousand words, has yet to be tested.

6. SOME COMPARISONS AND CONTRASTS

In this section I shall compare and contrast, under some eight interconnected headings, the projects described in the body of the notes. This is not easy to do, particularly when the present author is among the writers discussed, though that is easily remedied by the reader's making an appropriate discount. A more serious problem is that, at this stage of research

in artificial intelligence and natural language, the most attractive distinctions dissolve on more detailed scrutiny, largely because of the lack of any precise theoretical statement in most, if not all, the major projects. There are those who think that it therefore follows that this is not the moment for any form of critical comparison in this field, and that no more is needed than a "positive attitude" towards all possible projects. Only those who feel that, on the contrary, any time is as good as any other for the discussion of intellectual differences in the hope of progress, should read on.

It must be admitted right away that the selection of projects discussed above, like Winograd's distinction between first and second generation systems, on which the selection was to some extent based, cannot be defended by any strict definitions: one that would, in this case, include all the projects described, and exclude all those of Winograd's "first generation". One might, for example, want to define second generation systems (in the study of natural alanguage within the AI paradigm) in some very general terms, such as those systems which, (1) contain complex semantic structures for the representation of text that are significantly different from the "surface structure" of the input, and (2) contain cognate structures representing conceptual and real world knowledge that is not explicitly present in the input text. Even so general a description of a "frames" type approach would not cover Charniak or Colby with the first point. Moreover, the second point would certainly cover Winograd's work, as well as other first generation approaches, so it is clear, at the outset of any comparison, that there is not even a simple and unequivocal definition which covers all and only the projects to be compared.

Level of Representation

One important line of current dispute among the second generation approaches concerns the appropriate level of representation for natural language. On the one hand are those like Tolby, and apparently Charniak, who hold that language can be in effect, self representing (or represented by what I called projections in Section I), while on the other hand, there are those like Schank and myself who hold that the appropriate level of computation for inferences about natulra language is in some reduced, or primitive, representation. I wrote "appears" in the case of Charniak because he holds that his structures are independent of any particular level of representation, or rather, that they could be realized at a number of levels of representation, depending on the subject area. However, there is no doubt that the representation in terms of predicates that he offers in his work appears to be in one-to-one correspondence with English words.

The strongest low-level approach is undoubtedly that of Colby, who straightforwardly faces the enormous mapping problems involved if the structures are at the English word level. It is important to realize that this dispute is ultimately one of degree, since no one would claim that every locution recognized by an intelligent analyser must be mapped into a "deep" representation. To take an extreme case, any snstem that mapped "Good MOrning" into a deep semantic representation before deciding that the correct response was also "Good Morning" would be making a serious theoretical mistake.

However, the most serious argument for a non-superficial representation is not in terms of the avoidance of mapping

difficulties, but in terms of theoretical perspicuity of the primitive structures, and this argument is closely tied to the defence of semantic primitives in general, which is a large subject not to be undertaken here. One of the troubles about semantic primitives is that they are open to <u>bad</u> defenses, which decrease rather than increase their plausibility. For example, some users of them for linguistic representation have declared them to have some sort of objective existence, and have implied that there is a "right set" of primitives open to empirical discovery. On that view the essentially linguistic character of structures of primitives is lost, for they then might as well be strings of binary numbers, or something equally opaque and non-linguistic. No great deal of thought is required to see that that simply could not be the case. What is the case is that there is a considerable amount of psychological evidence that people are able to recall either the actual words or the syntactic structure used. There is large literature on this subject, from which two sample references would be (Wettler '73) and (Johnson-Laird '74).

These results are, of course, no proof of the existence of semantic primitives, but they are undoubtedly supporting evidence of their plausibility, as is, on a different plane, the result from the encoding of the whole Webster's Third International Dictionary at Systems Development Corporation, where it was found that a rank-ordered frequency count of the words used to define other words in that vast dictionary was a list (omitting "the" and "a") which corresponded almost item-for-item to a plausible list of semantic primitives, derived a priori, by those actually concerned to code the

structure of word and sentence meanings.

It is important to distinguish the dispute about level from the, closely connected, topic that I shall call the centrality of the knowledge required by a language understanding system.

Centrality

What I am calling the centrality of certain kinds of information concerns not its level of representation but its non-specificity: again a contrast can be drawn between the sorts of information required by Charniak's system, on the one hand, and that required by Schank's and my own on the other. Charniyk's examples suggest that the fundamental form of information is highly specific* to particular situations, like parties and the giving of presents, while the sorts of information central to Schank's and my own systems are partial assertions about human wants, expectations, and so on, many of which are so general as to be almost vacuous which, one might argue, is why their role in understanding has been ignored for so long.

If I were a reasonably fluent speaker of, say, German, I m ight well not understand a German conversation about birthday presents unless I had detailed factual information about how Germans organize the giving of presents, which might be considerably different from the way we do it. Conversely, of course, I might understand much of a technical article about a subject in which I was an expert, even though I knew very little of the language in which it was written.

*In Charniak's most recent paper (1974), he gives much more general rules, such as his "rule of significant sub-action".

There are certainly considerations that tell for Charniak's approach, and it is perhaps a paradox that the sort of natural language understander that would tend to confirm his assumptions would be one concerned with discourse about, say, the details of repairing a motor car, where factual information is what is central, yet, ironically, Charniak has concentrated on something as general as children's stories, with their need of deep assumptions about human desires and behaviour.

In the end this difference may again turn out to be one of emphasis, and of what is most appropriate to different subject areas, though there may be a very general issue lurking somewhere here. It seems to me not a foolish question to ask whether much of what appears to be <u>about natural language</u> in AI research is in fact about language at all, Even if it is not that may in no way detract from its value. Newell (Moore, Newell '73) has argued that A.I. work is in fact "theoretical psychology", in which case it could hardly be research <u>on</u> natural language. When describing Winograd's work earlier in the paper, I raised this question in a weak form by asking whether his definition of "pickup" had anything to do with the natural language use of the word, or whether it was rather a description of how his system picked something up, a quite different matter.

Suppose we generalize this query somewhat, by asking the apparently absurd question of what would be <u>wrong</u> with calling, say, Charniak's work an essay on the 'Socio-Economic Behaviour of American Children Under Stress?' In the case of Charniak's work this is a facetious question, asked only in order to make a point, but with an increasing

number of systems in A.I. being designed not essentially to do research on natural language, but in order to have a natural language "front end" to a system that is essentially intended to predict chemical spectra, or play snakes and ladders or whatever, the question becomes a serious one. It seems to be a good time to ask whether we should expect advance in understanding natural language from those tackling the problems head on, or those concerned to build a "front end". It is clearly the case that any piece of knowledge whatever could be essential to the understanding of some story. The question is, does it follow that the specification, organization and formalization of that knowledge is the study of language, because if it is then all human enquiry from physics and history to medicine is a linguistic enterprise. And, of course, that possibility has actually been entertained within certain strains of modern philosophy.

However, I am not trying here to breathe fresh life into a philosophical distinction, between being about language and not being about language, but rather introducing a practical distinction, (which is also a consideration in favour of opting, as I have, to work on very general and central areas of knowledge) between specific knowledge, and central knowledge without which a system could not be said to understand the language at all. For example, I might know nothing of the arrangement of American birthday parties, but could not be accused of not understanding English even though I failed to understand some particular children's story. Yet, if I did not have available some very general partial inference such as the one people being hurt and falling, or one about people endeavouring to possess things that they want, then it is quite possible

that my lack of understanding of quite simple sentences would cause observers to think that I did not understand English. An interesting and difficult question that then arises is whether those who concentrate on central and less central areas of discourse could, in principle, weld their bodies of inferences together in such a way as to create a wider system; whether, to put the matter another way, natural language is a whole that can be built up from parts?

Phenomenological level

Another distinction that can be confused with the central-specific one is that of the "phenomenological levels" of inferences in an understanding system. I mean nothing daunting by the phrase: consider the action eating which is, as a matter of anatomical fact, quite often an act of bringing the bones of my ulna and radius (in my arm) close to that of my lower mandible (my jaw). Yet clearly, any system of common sense inferences that considered such a truth when reasoning about eating would be making a mistake. One might say that the phenomenological level of the analysis was wrong even though all the inferences it made were true ones. The same would be true of any A.I. system that made everyday inferences about physical objects by considering their quantum structure.

Schank's analysis of eating contains the information that it is done by moving the hands to the mouth, and it might be argued that even this is going too far from the "meaning" of eating, whatever that may be, towards generally true information about the act which, if always inferred about all acts of eating, will carry the system unmanageably far.

There is no denying that this sort of information might be useful to have around somewhere; that, in Minsky's terms, the "default" value of the instrument for eating is the hand brought to the mouth, so that, if we have no contrary information, then that is the way to assume that any given act of eating was performed. Nonetheless, there clearly is a danger, and that is all I am drawing attention to here, of taking inferences to a phenomenological level beyond that of common sense. A clearer case, in my view, would be Schank's analysis (1974a) of mental activity in which all actions, such as kicking a ball, say, are preceded by a mental action of conceiving or deciding to kick a ball. This is clearly a level of analysis untrue to common sense, and which can have only harmful effects in a system intended to mimic common sense reasoning and understanding.

<u>Decoupling</u>

Another general issue in dispute concerns what I shall call <u>decoupling</u>, which is whether or not the <u>actual parsing</u> of text or dialogue into an "understanding system" is essential. Charniak and Minsky believe that this initial "parsing" can be effectively decoupled from the interesting inferential work and simply assumed. But, in my view, that is not so, because many of the later inferences would actually have to be done already, in order to have achieved the initial parsing. For example, in analysing "he shot her with a colt", we cannot ascribe any structure at all until we can make the inference that guns rather than horses are instruments for shooting, and so such a sentence cannot be inserted into an inference-but-no-parsing structure (as in (AT JOHN STATION)) without assuming that language does not have one of its

essential characteristics, namely systematic ambiguity. As I argued at the beginning of the notes, if one allows representational structures to have significance quite independent of their application, then one may not be in a situation essentially different from that of the logician who simply asserts that such-and-such is the "right structure" of some sentence.

Also, the inferences required to resolve word sense ambiguities, and those required to resolve pronoun reference problems, are not of different types; often the two problems occur in a simple sentence and must be resolved together. But Charniak's decoupling has the effect of completely separating these two closely related linguistic phenomena in what seems to me an unrealistic manner. His system does inferencing to resolve promoun ambiguities, while sense ambiguity is presumably to be done in the future by some other, untimately recoupled, system.*

Another way of pointing up the difference between the attitudes of second generation systems to decoupling, in relation to the first generation, is by describing the role of syntax analysis in them. As we saw, syntax was the heart of Winograd's system, but both levels of frame approach discount syntax analysis, though for very different reasons: Charniak does so because it is part of the initial parsing from which his inferential work has been decoupled. Schank and I do so because we believe semantic analysis to be fundamental, and that in an actual implementation the results of syntactic analysis can all be achieved by a sufficiently powerful semantic analyser. And this last assumption is

*Although Charniak would argue that sense ambiguity could be introduced into his system in its present form.

confirmed by the limited degree of success that the two semantic analysers have actually achieved in operation.

Availability of surface structure

An issue close to that of the appropriate level of representation in a system is that of the availability of the surface structure of the language analysed; or, to put it more crudely, the availability during subsequent analysis of the actual words being analysed. These are clearly available in Colby, and are indirectly available in Winograd's and my own system, but Shcank makes a point of the importance of their non-availability, on the grounds that an ideal representation should be totally independent of the input surface structure and words. There are both theoretical and practical aspects to this claim of Schank's: in the limit, the order of the sentences of a text is part of its surface structure, and presumably it is not intended to abandon this "superficial information". In one of his recent papers (1974b) Schank seems to have accepted some limitation on the abandonment of surface structure.

The other, practical point concerns the form of representation employed: in the (1973) implementation of Schank's system using an analyser of input text, a memory and a generator of responses, it was intended that nothing should be transferred from the input program to the output program except a representation coded in the structures of primitives discussed earlier*. The question that arises is, can that structure specify and distinguish word-senses adequately without transferring information specifically associated with the input word? Schank clearly believes the answer to

* This point is to some extent hypothetical since, as we saw, Schank's conceptualizations still do contain, or appear to contain, many surface items; in particular nouns, adjectives and adverbs, However, this is a transitional matter and they are in the course of replacement, as noted, by non-superficial items.

this question is yes, but that cannot be considered established by the scale of computations yet described in print.

A suitable environment in which to consider the question is that of translation from one language to another: suppose we are analysing a sentence containing the word "nail", meaning a physical object. It is clear that the translation of that word into French should not be the same as the translation for "screw" or "peg". Yet is it plausible that any discription of the function of these three entities entirely in terms of semantic primitives, and without any explicit mention of the word name and its connection to its French equivalent, will be sufficient to ensure that only the right match is made?

Application

This point is a generalisation of the last two, and concerns the way in which different systems display, in the structures they manipulate, the actual procedures of application of those structures to input text or dialogue. This is a matter different from both that of the availability of the surface structure, and of a computer implementation of the system. In the case of Colby's patterns, for example the form of their application to the input English is clear, even though the matching involved could be achieved by many different implementation algorithms. In the case of my own system, I hold the same to be true of the template structures, even though by the time the input has reached the canonical template form it is considerably different from the input surface structure. The system at the extreme end of any scale of perspicuity of application is Winograd's, where the procedural notation, by its nature, tends to make clear

the way in which the structures are applied. AT the other end are the systems of Schank and Charniak, where no application is specified, which means that the representations are not only compatible with many implementation algorithms, which does not matter, but are also compatible with many systems of linguistic rules, whose specification is an essential piece of inquiry, and whose subsequent production may cause the basic system to be fundamentally different.

English prepositions will serve as an example: in Schank's case notation there is no indication of how the case discriminations are actually to be applied to English prepositions in text. So, for example, the preposition "in" can correspond to the containment case, time location, and spatial location, among others. As we saw earlier, the discrimination involved in actual analysis is a matter of specifying very delicate semantic rules ranging over the basic semantic structures employed. Indeed, the structures and case system themselves seem to me to be essentially dependent on the nature and applicability of such rules*, and so this application of the system should have an obvious place in the overall structures. It is not seomthing to be delegated to a mere "implementation". If enough of the linguistic intractables ** of English analysis

---

* This is not meant to be just bland assertion. I have written at some length on the relations between application and the theoretical status of linguistic theories in (Wilks '74)

** The differences between Minsky's (1974) notion of "default value" and what I have called "preference" can be pointed up in terms of application. Minsky suggest "gun" as the default value of the instrument of the action of shooting, but I would claim that, in an example like the earlier "He shot her with a colt", we need to be able to see in the structure assigned wiether, or not what is offered as the apparent instrument is in fact an instrument and whether it is the default or not. In other words, we need sufficient structure of application to see not only that "shooting" prefers an instrument that is gun, but also why it will choose the sense of "colt" that is a gun rather than the one which is a horse.

were to be delegated out of the representation, A.I. would be offering no more to the analysis of natural language than the logicians who proffer the predicate calculus as a plausible structure for English.

In some of his more recent writings Winograd has begun to develop a view that is considerably stronger than this 'application'one: in his view the control structure of an understanding program is itself of theoretical significance, for only in that way, he believes, can natural language programs of great size and complexity remain perspicuous.

Forward inference

Another outstanding dispute concerns whether one should make massive forward inferences as one goes through a text, keeping all one's expectations intact, as Charniak and Schank hold, or whether, as I hold, one should adopt some 'laziness hypothesis' about understanding, and generate deeper inferences only when the system is unable to solve, say, a referential problem by more superficial methods. Or, in computer terms, should a system be problem or date-driven?

Although Shank sometimes writes of a system making "all possible" inferences as it proceeds through a text, this is not in fact the heart of the dispute, since no one would want to defend any strong definition of the term "all possible inferences". Charniak's argument is that, unless certain forward inferences are made during an analysis of, say, a story forward inferences, that is, that are not problem-driven; not made in response to any particular problem of analysis then known to the system then, as a matter of empirical fact, the system will not in general be able to solve ambiguity or reference problems that arise later, because it will never in fact be possible to locate (while looking backwards

at the text, as it were) the points where those forward inferences ought to have been made. This is, in very crude summary, Charniak's case against a purely problem-driven inferencer in a natural language understander.

A difficulty with this argument is the location of an example of text that confirms the point in a non-contentious manner. Charniak has found an excerpt from a book describing the life of apes in which it is indeed hard to locate the reference of a particular pronoun in a given passage. Charniak's case is that it is only possible to do so if one has made certain (non-problem occasioned) inferences earlier in the story. But a number of readers find it quite hard to refer that particular promoun anyway, which might siggest that the text was simply badly written.

This is a difficult matter about which to be precise: It would be possible, for example, to agree with Charniak's argument and still construct a purely problem-driven inferencer on the ground that, at the moment, this is the only way one can cope with the vast majority of inferences for understanding, since any system of inferences made in response to no particular problem in the text is too hard to control in practice. Indeed, it is noticeable that the most recent papers of Schank (1974a and 1974b) and Charniak (1974) have been considerably less forward-inference oriented than earlier ones*.

This dispute is perhaps only one of degree, and about the possibility of defining a degree of forward inference that aids the solution of later semantic problems without going into unnecessary depth**. This might be an area where psychological investigations would be of enormous help to workers in A.I.

* A particularly interesting withdrawl of a strong forward inference thesis is hidden away on p. 283 of (Rieger '74) but has been located by the keen eye of E. Charniak

** which may be no more than a psychological restatement of what used to be called (Hayes '71) (Sandewall '72) the frame problem. (no relation, P.E.)

The justification of systems

Finally, one might usefully, though briefly, contrast the different modes of justification implicitly appealed to by the systems described earlier in this paper. These seem to me to reduce to four:

(i) In terms of the power of the inferential system employed. This form of justification has underlain the early predicate calculus-based language programs, and is behind Hayes' (1974) recent demand that any formalism for natural language analysis should admit of a set theoretic semantics, in the Tarskian sense, so as to gain "intellectual respectability", as he puts it. The same general type of justification is appealed to in some degree by systems with PLANNER-type formalisms.

(ii) In terms of the provision and formalisation, in any terms including English, of the sorts of knowledge required to understand areas of discourse.

(iii) In terms of the actual performance of a system, implemented on a computer, at a task agreed to demonstrate understanding.

(iv) In terms of the linguistic and-or psychological plausibility of the proffered system of representation. Oversimplifying considerably, one might say that Charniak's system appeals mostly to (ii) and somewhat to (i) and (iv); Winograd's to (iii) and somewhat to the other three categories; Colby's (as regards its natural language, rather than psychiatric, aspects) appeals almost entirely to (iii); and Schank's and my own to differing mixtures of (ii), (iii) and (iv).

In the end, of course, only (iii) counts for empiricists, but there is considerable difficulty in getting all parties to

agree to the terms of a test. A cynic might say that, in the end, all these systems analyse the sentences that they analyse or, to put the same point a little more theoretically, there is a sense in which systems, those described here and those elsewhere, each define a natural language, namely the one to which it applies. The difficult question is the extent to which those many and small natural languages resemble English.

7. CONCLUSION

The last section stressed areas of current disagreement, but there would, if votes were taken, be considerable agreement among A.I. workers on natural language about where the large problems of the immediate future are: the need for a good memory model has been stressed by Schank (1974a), and many would add the need for an extended procedural theory of texts, rather than of individual example sentences, and for a more sophisticated theory of reasons, causes, and motives for use in a theory of understanding. Many might also be persuaded to agree on the need to steer between the Scylla of trivial first generation implementations and the Charybdis of utterly fantastic ones. By the latter, I mean projects that have been seriously discussed, but never implemented for obvious reasons, that would, say, enable a dialogue program to discuss whether or not a participant in a given story "felt guilty", and if so why.

The last disease has sometimes had as a major symptom an extensive use of the word "pragmatics" (though this can also indicate quite benign conditions in other cases), along with the implicit claim that "semantics has been solved, so we should get on with the pragmatics". It still needs repeating that there is no sense whatever in which the semantics of

natural language has been solved. It is still the enormous barrier it has always been, even if a few dents in its surface are beginning to appear here and there. Even if we stick to the simplest examples, that present no difficulty to the human reader---and it must be admitted that it has been one of the persistent faults of the A. I. paradigm of language that it has spent too much time on puzzles examples---there are still great difficulties both systematic and linguistic.

An example of the former would be the development of a synamic system of understanding texts or stories that had any capacity to recover after having its expectations satisfied and then, subsequently, frustrated.. At present no system of the sort described, whether of demons, preference or whatever, has any such capacity to recover. The situation is quite different from that in a dialogue, as in Winograd's system, where, on being given each new piece of information, the system checks it against what it knows, to see if it is being contradicted, and then behaves in an appropriately puzzled way if it is. In frame or "expectation" systems it is all too easy to construct apparently trick, but basically plausible, examples that satisfy what was being looked for and then overturn it. That possibility is already built into the notion of frame or expectation. An example of Phil Hayes against my own system will serve: consider "The hunter licked his gun all over, and the stock tasted especially good". What is meant by "stock" is clearly the stock piece of the gun, but any preference system like mine that considers the two senses of "stock", and sees that an edible, soup, sense of "stock" is the preferred object of the action "taste", will infallibly opt for the wrong sense. Any frame or expectation system is prone to the same general kind of counter-example.

In particular cases like this it is easy to suggest what might be done: here we might suggest a preference attached to the formula for anything that was essentially part of another thing (stock = 'part of gun" in this case), so that a local search was made whenever the "part-of" entity was mentioned, and the satisfaction of that search would always be the overriding preference. But that is not the same as a general solution to the problem, which used to be called that of "topic" in the computational semantics of the Fifties. There are no solutions to this problem available here and now, though some suggestions have been made by Abelson (1974) and McDermott (1974).

A closely related, but equally intractable, problem is that of how to combine highly specific factual information within a general semantic structuring. Systems like Charniak's are, as we saw, concerned with specific rather than conceptual information, but there are quite simple "semantic specificity" problems and one could not reasonably expect to be tackled even in a system devoted to the handling of facts, as can be seen by contrasting the sentences:

The deer came out of the wood.

The grub came out of the wood.

where we might safely assume that readers would assign quite different senses to "wood" in the two cases simply on the basis of the two different agents. No-one, to my knowledge, has suggested any general method for tackling such elementary examples.

But, to finish on the bright side, it is important to stress that there is indeed an A.I. paradigm of language understanding in existence*, one that embraces first and

---

* One of the very few acknowledgements of this fact, of the possibility of an A.I. paradigm of language, from a linguist is (Fillmore '74).

second generation approaches, and which goes back, I suggested, to a considerable amount of earlier work in computational linguistics. It can be distinguished by a catalogue of neglect by conventional linguistics that can be summarised under three heads:

(i) Theories of language must have procedural application to the subject matter that could in principle result in computer application and subsequent empirical test.

(ii) theories of language must deal with it in a communicative context, one amenable to empirical assessment. Merely sorting, as generative theories were designed to do, is not enough.

(iii) theories of language must also be, in clear sense, theories of the formalisation and organization of knowledge. If they are not then we can know in advance that they can never tackle the problem of language understanding.

A NOTE ABOUT THE REFERENCES:

The references are given only for completeness. It is not of course expected that those attending the course should be familiar with them all. However, it would be good to have glanced at the details of an article by each of the authors surveyed. In the abbreviated form of references below: MOD and AIJC3 each contain articles by 3 of the 4 authors (not the same 4).

In order to compress the reference list the following abbreviations for collections of articles will be used.

AIJC3 Advanced papers of the Third International Joint Conference on artificial Intelligence, Stanford, Calif. 1973

AISB Proceedings of the Summer Conference of the Society for Artificial Intelligence and Simulation of Behaviour, University of Sussex. 1974

TEDD Proceedings of the First International Conference on Machine Translation, National Physical Laboratory, Teddington, 1961 (HMSO, London, 1961).

ACL Proceedings of the Conference of the Association for Computational Linguistics, Amherst, Mass. 1974.

CAST Memoranda from the Istituto per gli studi Semantici e Cognitivi, Castagnola, Switzerland

MITAI Memoranda from the Artificial Intelligence Laboratory, Massachusetts Institute of Technology.

SUAIM Memoranda from the Artificial Intelligence Laboratory, Stanford University, Stanford, Calif.

SRITN Technical notes from the Stanford Research Institute, Menlo Park, Calif.

MOD Papers in Computer MOdels of Thought and Language, ed. by Schank and Colby, (Freeman, San Francisco, 1973).

(Abelson '73)
Abelson, R. P., The Structure of Belief Systems, in MOD.

(Abelson '74)
Abelson, R. P., Discourse Headers, to appear in Papers in Memory of J. Carbonell.

(Colby, Parkison '74)
Colby, K. M., Parkison, R. C., Pattern-Matching Rules for the Recognition of Natural Language Dialogue Expressions, in American Journal of Computational Linguistics, Vol. 1, No. 1.

(Colby, Hilf '73)
Colby, K. M., Hilf, F. D., Multidimensional Analysis in Evaluating a Simulation of Paranoid Thought, in SUAIM-194.

(Colby et al '71)
Colby, K. M., et al, Artificial Paranoia, Artificial Intelligence 2.

(Enea, Colby '73)
Enea, H., Colby, K. M., Idiolectic Language-Analysis for Understanding Doctor-Patient Dialogues, AIJC3.

(Fillmore '74)
Fillmore, C., Summer School lectures, University of Pisa, Italy.

(Goldman '74)
Goldman, N. M., Cumputer Generation of Natural Language from a Deep Conceptual Base, CAST No. 2.

(Halliday '70)
Halliday, M. A. K., in Foundations of Language.

(Hayes '71)
Hayes, P. J., The Frame Problem and Related Problems in Artificial Intelligence, in SUAIM-53

(Hayes '74)
Hayes, P. J., Some Problems and Non-Problems in Representation Theory, in AISB

(Hays '64)
Hays, D., Dependency Theory: A Formalism and Some Observations, Language, Vol. 40.

(Herskovits '73)
Herskovits, A., The generation of French from a semantic representation, SUAIM-212.

(Johnson-Laird '74)
Johnson-Laird, P., Memory for words, Nature, Vol. 251.

(McDermott '74)
McDermott, D. V., Assimilation of New Information by a Natural Language Understanding System, MITAI, Tr-291.

(Minsky '74)
Minsky, M., Frame systems, unpublished MSS from Massachusetts Institute of Technology.

(Rieger '74)
Rieger, C. J., Conceptual Memory, Stanford University Ph.D. Thesis.

(Riesbeck '74)
Riesbeck, C. K., Computational Understanding: Analysis Sentences and Context, CAST No. 4.

(Sandewall '72)
Sandewall, E1, PCF-2, A First-Order Calculus for Expressing Conceptual Information, Uppsala University, Department of Computer Science.

(Schank '73a)
Schank, R. C., The Fourteen Primitive Actions and Their Inferences, SUAIM-183.

(Schank '73b)
Schank, R. C., Identification of Conceptualizations Underlying Natural Language, in MOD

(Schank '74a)
Schank, R. C., Causality and Reasoning, CAST NO. 1.

(Schank '74b)
Schank, R. C., Understanding Paragraphs, CAST No. 6.

(Schank et al. '70)
Schank, R. C., et al., SPINOZA II: Conceptual Case-Based Natural Language Analysis, SUAIM- 109.

(Schank, et al '73)
Schank, R. C., et al, MARGIE: Memory, Analysis, Response Generation and Inference on English, in AIJC3.

(Wettler '73)
Wettler, M., Zur Speicherung Syntaktischer Merkmale in Langzeitgedachtnis, Bericht über den 28, Konress der Deutschen Gesellschaft f. Psychologie.

(Wilks '73a)
Wilks, Y. A., Preference Semantics, in SUAIM-206, and to appear in Keenan (ed.) Formal Semantics of Natural Language, (Cambridge University Press. 1975)

(Wilks '73b)
Wilks, Y. A., An Artificial Intelligence Approach to Machine Translation, in MOD.

(Wilks '74)
Wilks, Y. A., One Small Head--Models and Theories in Linguistics, in Foundations of Language, Vol. 11.

(Wilks in press)
Wilks, Y. A., A preferential, Pattern Seeking, Semantics for Natural Language Inference, to appear in Artificial Intelligence.

(Weber '72)
Weber, S., Semantic Categories of Nominals for Conceptual Dependency Analysis of Natural Language, SUAIM-172

(Winograd '71)
Winograd, T., Procedures as a Representation for Data in a Computer Program for Understanding Natural Language, MIT Ph.D. Thesis.

(Winograd '72)
Winograd, T., Understanding Natural Language, (Academic Press, Inc., New York, 1972).

(Winograd '73)
Winograd, T., A Survey Lecture on Natural Language, Third International Joint Conference on Artificial Intelligence, Stanford, Ca. 1973.

(Wittgenstein '53)
Wittgenstein, L., Philosophical Investigations, (Blackwell, Oxford, 1953).